

Cisco

MPLS Transport Profile (MPLS-TP) Solution

Protocol Architecture and Functionality Verification Test Report

Introduction

In February 2011, Cisco commissioned EANTC with an in-depth review of Cisco's new MPLS-TP implementation across three switching platforms and a management solution.

We took a look at how Cisco answers the main questions related to MPLS-TP: How are connection-oriented Ethernet services configured and provisioned? How are services monitored and managed? What is the maximum outage expected in the case of a failure? What Operation, Administration and Maintenance (OAM) tools are available for localizing failures? This report outlines the results of our rigorous tests which Cisco's products undertook.

EANTC Validation Highlights

- **Unified service configuration across MPLS-TP and IP/MPLS using existing pseudowire technology**
- **Comprehensive fault management including tools for fault verification and isolation**
- **1:1 linear protection with hitless restoration and hitless manual switchover with sub-second protection switching**
- **Integrated monitoring of packet transport technology using Cisco Prime Network**
- **Standards based solution**

Tested Devices and Architecture

Cisco constructed the test bed to represent a carrier network with an IP/MPLS core and an MPLS-TP transport/aggregation network area. Two test scenarios were defined, each with a different logical aggregation topology. Cisco brought to the test the 7604, 7606, ASR 9006, ASR 9010, and CPT 600. A single physical topology was used for both logical scenarios for simplicity. All devices used pre-release software. No rebooting or software upgrades took place on any Cisco equipment during the test.

We ran through each test twice — once for each of the two logical scenarios configured by Cisco. Both scenarios included an IP/MPLS domain and an MPLS-TP domain — the difference being the logical roles by each device. Please see figure 2 for details. In the first scenario, for example, the Cisco ASR 9010 interconnected the two domains while the Cisco 7604 played this role in the second scenario. In both cases the Cisco ASR 9006 was the only other device in the IP/MPLS domain, the rest being in the MPLS-TP domain.

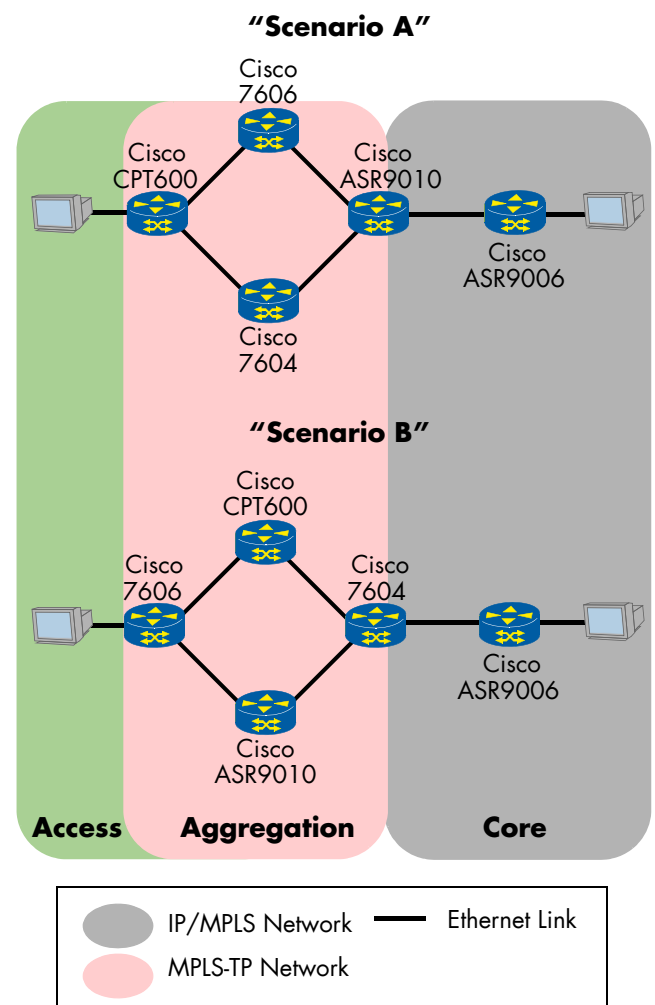


Figure 1: Test Setup and Logical Scenarios

The two domains (MPLS-TP and IP/MPLS) had the following differences:

1. No signaling protocols were used in the MPLS-TP domain. All Pseudowires and their Label Switched Paths (LSPs) were entered manually into the Forwarding Information Base (FIB) by configuring incoming and outgoing labels and interfaces in each case via the respective device's command line interface (CLI). The IP/MPLS domain used OSPF-TE and RSVP-TE for computation of label switched paths (LSP) and the signaling of labels. In addition, it used LDP for pseudowire signaling.
2. The MPLS-TP domain implemented a set of new OAM protocols and features being discussed and defined at the IETF. (e.g. continuity check, connectivity verification, route tracing, fault management). The IP/MPLS network used the OAM mechanism for fault verification and isolation using LSP Ping/Trace which have been defined for some years now.

EANTC verified that the appropriate configuration was found on each device, that such labels and forwarding was used for both data plane traffic and OAM, and ensured that the tester emulating customer traffic had no IP configured at all. Interface identifiers were configured together with subnet masks as defined in draft-ietf-mpls-tp-identifiers.

Cisco Prime™ Network

Throughout the tests, Cisco demonstrated their management system - Cisco Prime Network. The system provided monitoring and assurance of packet and transport devices from a device view (physical and logical) up to a topology and service view. The application monitored all devices under test using SNMP and automated command-line inference (CLI). The software run on a server connected to the management network.

To simplify operation and facilitate monitoring, two main views were used throughout the testing:

1. Cisco Prime Network topology view, where the device and topological link were reported.
2. Cisco Prime Network service view, where the user can drill down on the information of a selected service (pseudowire, MPLS-TP tunnel, etc.)

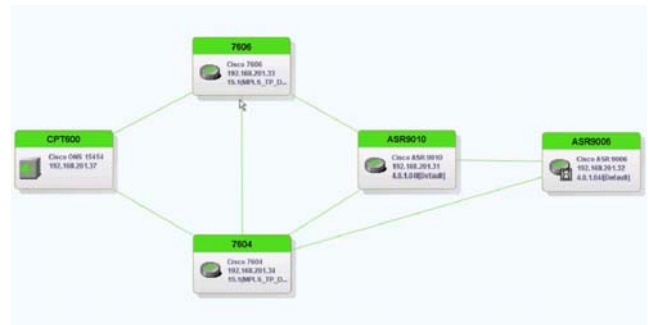


Figure 2: Cisco Prime Network - Topology View

Cisco Prime Network reported and summarized all alarms in both the topological as well as in the service view. Moreover, the service view automatically updated informing the user of the current alarmed devices as well as of the path switching when happening.

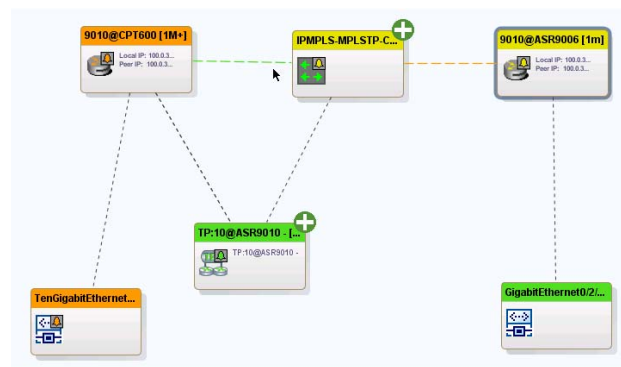


Figure 3: Cisco Prime Network - Service View

Cisco Prime Network was able to detect and display changes in the network services thereby enabling rapid fault isolation and restoration of services. For MPLS-TP services these management capabilities are critical for delivering high-quality services.

Operation, Administration and Maintenance (OAM) Tests

An important component of a transport profile for MPLS consists of defining OAM tools that will enable the same level of operational control, troubleshooting and Service Level Agreement (SLA) validation that service providers have in their legacy transport networks. As the requirements for OAM in MPLS transport profile IETF request for comments (RFC 5860) specifies these OAM mechanisms must operate in band, must be applicable to all services and must not rely on a control plane.

In all the OAM tests we performed, we verified that the OAM tools were using the Generic Associated Channel and Generic Associated Label (GACH/GAL) as described in RFC 5586.

MPLS-TP Trouble Shooting Tools – LSP Ping and Traceroute

LSP Ping and Traceroute provide the on-demand connectivity verification (CV) and path tracing functions as part of the MPLS-TP OAM toolkit. These tools facilitate the detection and isolation of unexpected connectivity (e.g. merging, mis-connection) across an MPLS-TP LSP.

LSP Ping and Traceroute have been standardized, implemented and deployed for some years in IP/MPLS networks. For MPLS-TP, new extensions have been standardized to support the use of the GACH, introduce support for static LSP and PW, and introduce a non-IP encapsulation option. These new extensions are defined in draft-ietf-mpls-tp-on-demand-cv.

In all test scenarios Cisco executed LSP ping and traceroute commands from both ends of the MPLS-TP domain using both the CLI, and Cisco Prime Network management solution. Both functions were executed on working and protection paths in active and standby state. In the case of an emulated link failure, as described later in this report, LSP Ping was re-executed to show that the tool was unable to reach the remote peer as planned, and LSP Traceroute was used to verify that the correct number of hops responded, and the Traceroute was interrupted where expected. LSP Ping and Traceroute commands were also executed within the IP/MPLS domain.

Fault Management

Fault Management defines new OAM extensions to automatically indicate a disruption on a link or node along the path to an MPLS-TP LSP end point. These indications are important to suppress alarms and to activate protection and are being defined as part of draft-ietf-mpls-tp-fault. In a traditional IP/MPLS network, the control plane typically signals such disruptions.

Cisco demonstrated two different fault management tools, Link Down Indication (LDI) and Locked Report (LKR). LDI is sent, as the name indicates, when a link within the MPLS-TP network is disrupted. On the other hand, LKR is generated when a link along the path is administratively shut down. Both are actions of a mid-point (or Label Switched Router / LSR) within the network, which signals the respective case to the end point which could not otherwise immediately notice

the failure before further alarms occur (alarm suppression).

In order to verify that the devices are able to correctly signal link down, we disrupted the link between the mid-point and the impairment device (sitting between the mid-point and the node interconnecting the two domains) and verified that a) the network switched over and made the backup path the active path, and b) the appropriate (LDI) status was shown on the MPLS-TP ingress. In the case of LKR, we administratively shut down the same link and similarly verified that a) the expected LSP switchover occurred, and b) the appropriate (LKR) status was displayed on the LSP end point.

Continuity Check (CC) and Remote Defect Indication (RDI)

Bidirectional Fault Detection (BFD), originally standardized across several RFCs as a protocol used to monitor IP network services bidirectionally at potentially high frequencies, was reused in the IETF to standardize the Continuity Check (CC) requirement to pro-actively monitor MPLS-TP LSPs. These extensions are specified under draft-ietf-mpls-tp-cc-cv-rdi.

Just as in other technologies with an RDI feature like SONET/SDH or Ethernet Connectivity Fault Management (CFM), MPLS-TP includes an RDI function so an LSP endpoint can signal a locally detected defect to the remote endpoint. For MPLS-TP, the existing BFD diagnostic field performs the RDI function as defined in draft-ietf-mpls-tp-cc-cv-rdi. We tested this function by blocking traffic from the node interconnecting the domains towards the downstream direction on the primary path. This unidirectional failure caused the node which stopped receiving BFD packets to indicate a local timer expired event in its sent BFD packet. The peer in turn switched over to the backup path, and reflected the RDI in the CLI. Additionally, an alarm was raised on Cisco Prime Network, which reflected the switchover as well.

Protection Switching

Linear protection is an important requirement for transport networks. Cisco demonstrated 1:1 protection with revertive behavior, manual switchover and multiple trigger mechanisms using BFD as a continuity check mechanism.

In our protection tests, Cisco configured redundant LSPs with unique physical paths, and enabled BFD to monitor each of them. We used an impairment generator to disrupt the forwarding plane without affecting

the physical link. This methodology, albeit more complex than traditional failure scenarios, ensured that the loss of BFD would be the reason for the network to detect an issue on the forwarding plane and switch traffic to the backup LSP. Three kinds of failure scenarios were emulated, unidirectional failures in each direction, and bidirectional failure. In all cases physical Ethernet link remained intact.

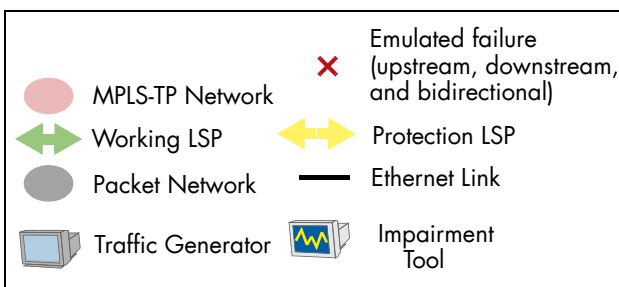



Figure 4: Tested Failure Scenarios

In order to measure the time taken for the network to recognize the failure, and switch traffic to the backup path, we sent bidirectional traffic using a traffic generator connected at each end of the network as shown in Figure 4. We used the amount of traffic loss observed to calculate the out of service time. For simplicity we used a rate of 1,000 packets per second, with a frame size of 128 bytes. This meant that the loss of 1 frame translated to 1 ms of service interruption.

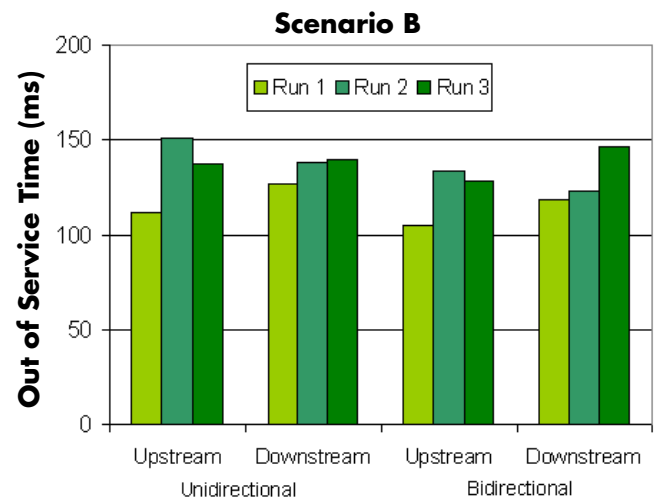
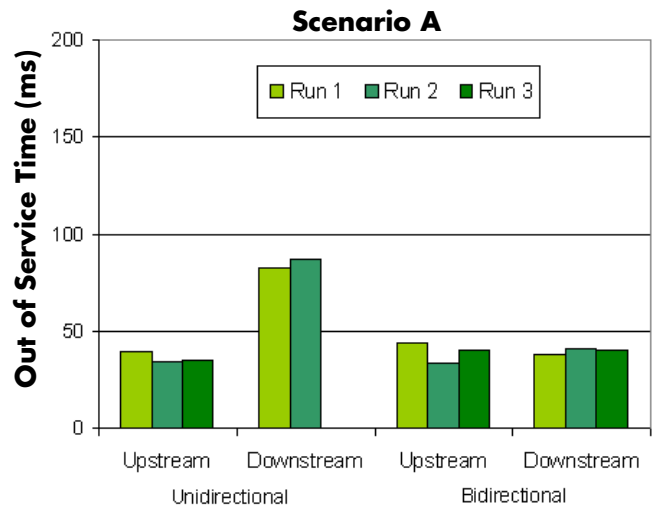


Figure 5: Link Protection Results¹

All traffic was VLAN tagged, with no IP header configured on the emulator. We expected the out of service time to be somewhere between two and three times the BFD CC interval, since all devices were configured to failover to the backup path once three Continuity Check (CC) messages were not received (the time-out was set to three missing packets). One would notice that the results between the two scenarios vary, but this is explained by the different speeds supported by the different devices. The ASR 9010 and the CPT 600

¹Figure 5 shows out of service times measured during failure events. All restoration events experienced no packet loss. In the case of Scenario B, Cisco later demonstrated an updated CC implementation on the 7604 which supported a message interval of 10 ms and resulted in sub-50 ms out of service times.

were configured with a CC interval of 15 milliseconds (ms) while the 7604 was configured with a CC interval of 50 ms.

Reversion was enabled for all protection tests as defined in draft-ietf-mpls-tp-survive-fwk. Traffic was redirected back to the original working LSP once the initial switchover trigger disappeared (e.g. fault affecting the working LSP). A wait-to-restore timer was configured such that the working LSP had to remain up for a period of time before the protection LSP transitioned to standby state and the working LSP transitioned to active state. This mechanism prevents excessive switchovers between the working and protection LSP that would result in additional traffic loss due to an intermittent fault. For all protection tests, no packet loss was experienced during reversion.

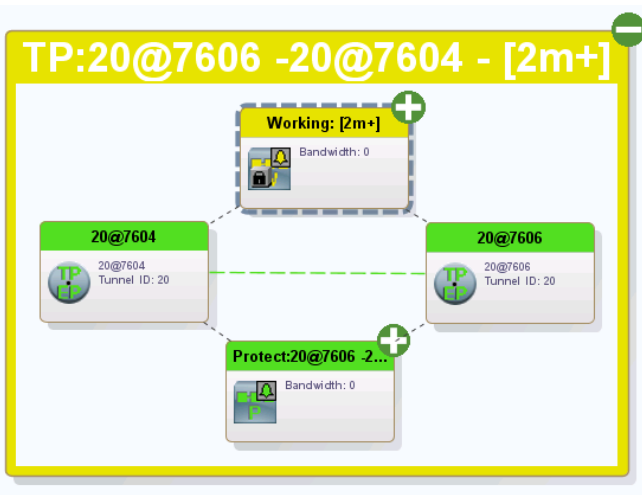


Figure 6: Cisco Prime Network - Reporting LSP Lockout

In addition, we verified manual switchover initiated from Cisco Prime Network. This switchover was initiated using a "lock-out" function on the working LSP from one end point. Cisco Prime Network displayed the working LSP as locked and the protection LSP as active after the switchover, with a color coding scheme. After the lock was removed, traffic was restored to the working LSP with the same revertive behavior described above. No packet loss was experienced on manual switchovers when the lock was introduced or removed.

End-to-end Pseudowire Status Notification

In IP/MPLS networks, LDP is typically used to signal pseudowires as defined in RFC 4447. End points can use LDP Notification messages to exchange status

information for pseudowires and the attachment circuits. When a pseudowire endpoint receives a "down" status notification from the remote end point, it can perform different actions including disabling the local attachment circuit, propagating the failure notification on the local attachment circuit, or switching to a standby pseudowire.

In our tests, Cisco demonstrated the use of the pseudowire status notification across the MPLS-TP domain in addition to the IP/MPLS domains. Due to the static nature of the pseudowire segment in the MPLS-TP network, the status notification was sent using the GACH as defined in draft-ietf-pwe3-static-pw-status. In the IP/MPLS network, the LDP status notification message was used. In both test scenarios, the Provider Edge (PE) nodes tore down the links connected to the traffic generators, which emulated the customer, and watch the multi-segment pseudowire come down one segment at a time on Cisco Prime Network. Figure 7 shows the topology in an orange state since the pseudowire was taken down (the green backup is still green) behind the window logging the specific alarms.

The screenshot shows the Cisco Prime Network interface with a topology view and an alarm log. The topology view shows a pseudowire between 9020@7606 and 9020@ASR9006. The alarm log shows the following entries:

Alarm Correlation	Short Description	Location	Acknowledged	Affected Devices Count
111475	Part down due to oper	7606#2:GigabitEthernet2/1	No	2
111477	SRMP Link down	7606#2:GigabitEthernet2/1	No	
111478	Line down syslog	7606#2:GigabitEthernet2/1	No	
111479	Link down syslog	7606#2:GigabitEthernet2/1	No	
111476	EFP oper down	7606#2:GigabitEthernet2/1 EFP-20	No	
111401	Layer 2 tunnel down	9020@7606<->9020@7604	No	
111480	Pseudo wire tunnel down	9020@7606	No	

Figure 7: Cisco Prime Network - Pseudowire Status Notification

Cisco demonstrated different alternatives in the use of the pseudowire control word. Operators can use this field to carry Layer-2 flags for ATM and Frame Relay, to enable specific pseudowire OAM modes, or to ensure in-order delivery of frames. In our tests, "Scenario A" did not use of the control word while "Scenario B" did. Pseudowire status notification was tested in both scenarios. The lack of use of the control word on scenario A did not have any negative impact on the other functions which were verified and described in this report.

Summary

Cisco's MPLS Transport Profile implementation, coupled with Cisco Prime Network demonstrated a robust solution aimed at service providers looking to move into Ethernet based connection-oriented transport solution. Using Cisco Prime Network we were able to visualize transport services and display failure events to using a graphical user interface. For MPLS-TP services these management capabilities are critical for delivering high-quality services.

We verified that the MPLS-TP tool-kit provided by Cisco is mature, standard based and robust. An operator used to transport network would be glad to find the tools he is used to seeing in Cisco's MPLS-TP solution.

For more information on the products tested:

Cisco ASR 9000

<http://www.cisco.com/en/US/products/ps9853/index.html>

Cisco 7600

<http://www.cisco.com/en/US/products/hw/routers/ps368/index.html>

Cisco CPT 600

<http://www.cisco.com/en/US/products/ps11348/index.html>

Cisco Prime Network

<http://www.cisco.com/go/prime>

About EANTC



The European Advanced Networking Test Center (EANTC) offers vendor-neutral network test services for manufacturers, service providers and enterprise customers. Primary business areas include interoperability, conformance and performance testing for IP, MPLS, Mobile Backhaul, VoIP,

Carrier Ethernet, Triple Play, and IP applications.

EANTC AG, Einsteinufer 17, 10587 Berlin, Germany
info@eantc.com, <http://www.eantc.com/>

V1.0 20110426